

Markov chain Monte Carlo searches for Galactic binaries in Mock LISA Data Challenge 1B data sets

Miquel Trias¹, Alberto Vecchio² and John Veitch²

¹ Departament de Física, Universitat de les Illes Balears, Cra. Valldemossa Km. 7.5, E-07122 Palma de Mallorca, Spain

² School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

E-mail: miquel.trias@uib.es, av@star.sr.bham.ac.uk, jveitch@star.sr.bham.ac.uk

Abstract. We are developing a Bayesian approach based on Markov chain Monte Carlo techniques to search for and extract information about white dwarf binary systems with the Laser Interferometer Space Antenna (LISA). Here we present results obtained by applying an initial implementation of this method to some of the data sets released in Round 1B of the Mock LISA Data Challenges. For Challenges 1B.1.1a and 1b the signals were recovered with parameters lying within the 95.5% posterior probability interval and the correlation between the true and recovered waveform is in excess of 99%. Results were not submitted for Challenge 1B.1.1c due to some convergence problems of the algorithm; despite this, the signal was detected in a search over a 2 mHz band.

PACS numbers: 04.80.Nn, 02.70.Uu, 07.05.Kf

1. Introduction

Galactic white dwarf (WD) binary systems are guaranteed sources for the future Laser Interferometer Space Antenna, LISA [1]. Despite the simple nature of the expected gravitational radiation – a quasi monochromatic signal – the data analysis task becomes challenging due to the tens of millions of such sources in the Galaxy that radiate in the instrument’s observational window, with signals strongly overlapping in time and frequency space at the LISA output [2, 3, 4]. As a consequence, LISA is expected to resolve about 10^4 galactic binaries with the remaining producing a “confusion noise”, whose exact magnitude will depend on the resolving power of the search methods (as well as the actual astrophysical population in the Galaxy).

Due to the large number of sources to be analysed at the same time and the fact that this number is unknown, Markov chain Monte Carlo (MCMC) techniques [5, 6, 7, 8, 9] and their extension to Reversible Jump Markov chain Monte Carlos (RJMCs) [10, 11, 12] are expected to be one of the most powerful search methods. There are several instances in which techniques based on Bayesian inference have been successfully implemented in the context of LISA data analysis, not only in the search for WD binary systems, but also gravitational radiation from massive-black-hole binary inspirals [13, 14, 15, 16, 17, 18] and extreme-mass ratio inspirals [19, 20].

In the summer of 2007, the Mock LISA Data Challenge (MLDC) Task Force [21, 22] released a re-issue of the Round 1 challenges, called Challenge 1B [23]. In

this paper, we present results obtained by applying an initial implementation of a Bayesian approach based on a MCMC method to the single galactic binary data sets, Challenge 1B.1.1a-c [24]. Several other groups have tackled this problem using a number of approaches: Cornish & Crowder [5, 6, 7, 8] developed algorithms based on variations of Monte Carlo Metropolis-Hastings Samplers and successfully applied them to searches for overlapping sources; MCMC methods for single-source analysis have been explored by several groups, *e.g.* [9]; Prix & Whelan [25, 26] and Królak and collaborators have developed a matched-filtering approach, based on the computation of the \mathcal{F} -statistic [27]. A summary of the results obtained in Challenge 1B is provided by the MLDC Task Force [23].

All the results that we present here, correspond to the blind challenge data sets and were obtained before the release of the key files, in December 2007.

2. Analysis method

Following a Bayesian approach, we can infer the probability density functions (PDFs) of the vector of the unknown model parameters $\vec{\lambda}$, given a data set d and some prior information W , using Bayes' theorem:

$$p(\vec{\lambda}|d, W) = \frac{p(\vec{\lambda}|W) p(d|\vec{\lambda}, W)}{p(d|W)}. \quad (1)$$

Here, $p(\vec{\lambda}|d, W)$ is the *posterior* probability density function of the parameters given the observed data (*i.e.* it is what we are interested in), $p(\vec{\lambda}|W)$ is the *prior* knowledge we have about the different parameters, $p(d|\vec{\lambda}, W)$ is the *likelihood function* of the data given the model and finally, $p(d|W)$ is a normalisation factor independent of the unknown model parameters, and therefore irrelevant to this MCMC analysis.

The gravitational wave (GW) signal emitted by a WD binary system with constant orbital frequency is characterised by 7 independent parameters: the frequency of the signal f , its amplitude \mathcal{A} , two angles to define the sky location of the source – here longitude ϕ and latitude ℓ – the GW polarisation ψ , the inclination angle between the angular momentum of the system and an unitary vector parallel to the line of sight, ι and finally, a constant value to fix the initial phase of the signal φ_0 . In our analysis we also treat the noise level that affects the measurements as unknown; we parametrise it with σ^2 the (constant) variance of the noise contribution in frequency domain in the small band where the signal lies. As a consequence, the analysis of the data containing a single galactic binary requires the estimation of an 8-dimensional parameter vector $\vec{\lambda}$.

MCMC methods are well suited to compute the joint posterior PDF $p(\vec{\lambda}|d, W)$, and the *marginalised* posterior PDF for any given parameter (or subset of parameters), say λ_1 :

$$p(\lambda_1|d, W) = \int d\lambda_2 \dots \int d\lambda_8 p(\vec{\lambda}|d, W). \quad (2)$$

From the PDF above, one can then compute the posterior mean as

$$\bar{\lambda}_j = \int_{-\infty}^{\infty} d\lambda_j \lambda_j p(\lambda_j|d, W). \quad (3)$$

We have implemented a Metropolis-Hastings MCMC algorithm which has the property that after an initial “burn-in” period (which is discarded in the generation

of the PDFs), it returns samples of $\vec{\lambda}$ with a probability density equal to the desired posterior $p(\vec{\lambda}|d, W)$ [28].

In an MCMC algorithm, a sequence of points ('a chain') is constructed. The first point is chosen randomly according to a uniform prior distribution, then subsequent elements of the chain are generated in the following way:

- (i) Given the current member of the chain corresponding to the parameter vector $\vec{\lambda}$, the new member $\vec{\lambda}'$ is proposed according to

$$\lambda'_i = \lambda_i + u \Delta_i \quad (i = 1, \dots, 8), \quad (4)$$

where u is a Gaussian random number with zero-mean and unit variance, and Δ_i sets the size of the step (see the next Section for more details).

- (ii) The new member is accepted with a probability computed according to the Metropolis-Hastings ratio:

$$\alpha_{\vec{\lambda}, \vec{\lambda}'} = \min \left(1, \frac{\pi(\vec{\lambda}')}{\pi(\vec{\lambda})} \right), \quad (5)$$

where $\pi(\vec{\lambda}) \propto p(\vec{\lambda}|W) p(d|\vec{\lambda}, W)$ is the so-called target distribution.

Since the signal from a galactic binary is nearly monochromatic at the LISA output, it is advantageous to work in the Fourier domain, and to consider only the small frequency band where the signal power is concentrated. In our implementation we analyse only the band $B_w = 1.5 \left[2 \left(5 + 4\pi f_0 \frac{R}{c} \sin \theta \right) f_m + |\dot{f}_0| T_{obs} \right]$ around f_0 , where $f_m = 1 \text{ yr}^{-1}$, f_0 is the signal frequency at a given reference time, \dot{f}_0 its time derivative (for the signal model adopted in MLDC-1B $\dot{f}_0 = 0$), and 1.5 is a safety factor. We therefore FFT the time series X , Y and Z of the unequal arm pseudo-Michelson outputs in which the MLDC data are distributed [24], and construct the two noise-orthogonal Time Domain Interferometry (TDI) outputs [29]:

$$A = \frac{2X - Y - Z}{3}, \quad E = \frac{Z - Y}{\sqrt{3}}. \quad (6)$$

A and E for a given choice of the source parameters are then computed directly in the Fourier domain using an approximation and software implementation provided by Cornish and Littenberg [12] in which the LISA response function for nearly monochromatic signals is separated in a *fast* and a *slow* part. The *fast* part is computed analytically in the frequency domain, while the *slow* one can be evaluated efficiently by sampling it at a much slower rate in the time domain and then doing an FFT. It takes approximately 10^{-2} s, depending on the frequency bandwidth, to generate a single signal in the relevant range, and in order to avoid spurious oscillations at the edges of the frequency band due to windowing, we generate the signal in a band 3/2 times wider than what is needed, B_w , and then select only the central part.

The likelihood function that needs to be computed at each step of the MCMC Metropolis-Hastings algorithm is given by

$$p(d|\vec{\lambda}, W) \propto \frac{1}{\sigma^N} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{\alpha=1}^2 \sum_{k=1}^N \left| \tilde{d}_{\alpha,k} - \tilde{h}_{\alpha,k}(\vec{\lambda}) \right|^2 \right\}, \quad (7)$$

where α labels the TDI channel – A and E –, $k = 1, \dots, N$ the frequency bin, \tilde{d} the data set and $\tilde{h}(\vec{\lambda})$ the predicted signal (model) in the Fourier domain.

Table 1. Information about the Mock LISA Data Challenges 1B.1.1X, corresponding to a single galactic binary signal. Brackets represent the prior range of the two parameters for the different sources. The data sets are approximately 1 yr long with a cadence of 15 seconds.

Challenge	SNR	frequency (mHz)
1B.1.1a	[10 , 20]	[0.9 , 1.1]
1B.1.1b	[10 , 20]	[2.9 , 3.1]
1B.1.1c	[10 , 20]	[9 , 11]

3. Results

As a test of the initial implementation of the analysis algorithm, we analysed the three data sets of Challenge 1B.1.1, each of which contained a single galactic binary signal buried in Gaussian and stationary instrument noise. In Table 1 we summarise the main properties of the three challenge data sets; the main difference amongst them is the actual value and prior range of the signal frequency. In particular, in Challenge 1B.1.1c the frequency prior region was 10 times wider than for Challenge 1B.1.1a-b.

The algorithm was exactly the same for each analysis and we adopted constant values for the amplitude of the proposals Δ_i throughout the evolution of the chains. For the angles, the value of Δ_i was set to one third of the prior range; for the frequency we used $\Delta_f = 3 \text{ yr}^{-1} = 95 \text{ nHz}$, corresponding to three frequency bins; we evolved the noise level in a logarithmic scale with $\Delta_{\log \sigma} = 0.01$, and for the signal’s amplitude we used $\Delta_A = 3 \times 10^{-25}$. These choices were motivated by reasons of efficiency, however for lack of time no specific tuning of the MCMC code took place.

Some additional care is required to probe efficiently the prior range of several source parameters. The signal model adopted for galactic binaries is characterised by 5 angular parameters, each of which with an intrinsic periodicity: initial phase φ_0 and longitude ϕ are modulo 2π ; latitude ℓ , polarisation angle ψ and orientation ι are modulo π . We have limited our parameter space in order to work always with angular values between 0 and their period. The co-latitude θ , that we use in our search instead of latitude, is taken between 0 and π , giving a latitude range in the usual interval $\ell \in (-\pi/2, \pi/2)$. There exists another symmetry concerning the angular parameters, and it is given by the fact that the transformations $\varphi_0 \rightarrow \varphi_0 \pm \pi$ or $\psi \rightarrow \psi \pm \pi/2$ produce a change of the waveform’s global sign. This implies that we can use *e.g.* $\varphi_0 \in (0, 2\pi)$ as before, but restrict the polarisation angle range to $\psi \in (0, \pi/2)$, since the $(\pi/2, \pi)$ range can be covered by adding or subtracting π to the initial phase.

The most crucial parameter for the search for galactic binaries is the frequency, and it is advantageous to analyse any given data set in a number of narrow frequency bands. For Challenge 1B.1.1a-b, the signals were so strong that their frequency could be approximately determined by simply computing the power spectrum of the data; the MCMC algorithm was then run on a narrow band around that frequency (see Section 3.1). Challenge 1B.1.1c was more challenging for us because the prior frequency range was 10 times larger and, the signal being at higher frequency, the Doppler effect spreads the power over many frequency bins without producing clearly identifiable peaks in the spectrum (see Section 3.2). The results of the analyses are discussed in some detail in the next two subsections.




Figure 1. Power spectral density of LISA's channel *A* for Challenge 1B.1.1a and 1b. The peak produced by the actual signal is highlighted by an ellipse.

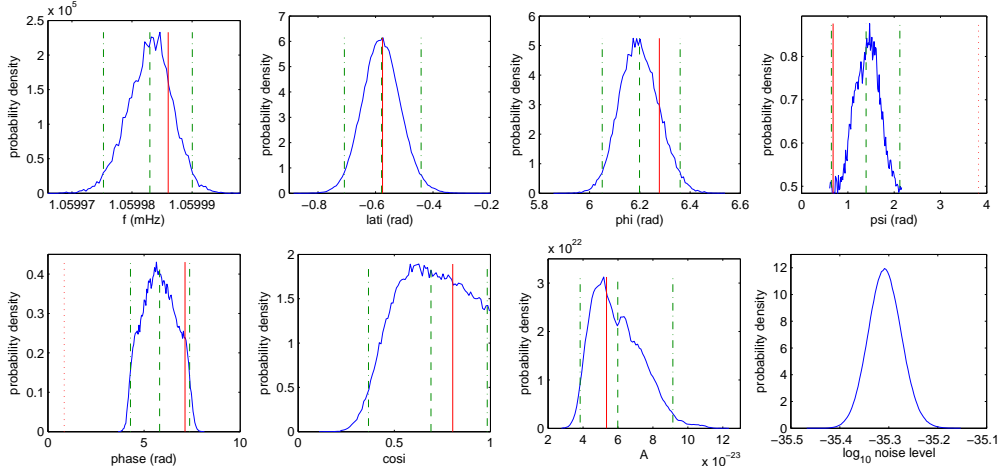


Figure 2. Marginalised posterior PDFs of each of the seven parameters that characterise the signal, and noise level for Challenge 1B.1.1a data set. The true values of the parameters are represented by a vertical solid line, while our submitted values (mean value of the distribution) together with their 95.5% probability interval [30] are represented by the dashed and dash-dot lines, respectively. Dotted lines indicate positions in the parameter space equivalent to the true ones.

3.1. Challenge 1B.1.1a-b

In Figure 1 we plot the power spectral density of the *A* TDI output from Challenge 1B.1.1a-b data sets in the relevant frequency range. It was easy to identify “by eye” the approximate frequency of the signal by simply looking at the highest peaks and checking for their presence in the noise orthogonal TDI channel, *E*. Using this procedure, we identified a narrow frequency range (50 frequency bins, corresponding to 1.58×10^{-6} Hz) for the analysis with the MCMC code to generate the posterior PDFs of the signal parameters. The typical length of the burn-in stage was chosen to be 2.5×10^4 steps, and the code run until 10^5 iterations of the Markov chain were completed, which in all cases took less than two hours to run on a 2.16 GHz processor. The results submitted for the challenge were obtained by combining the post-burn-in

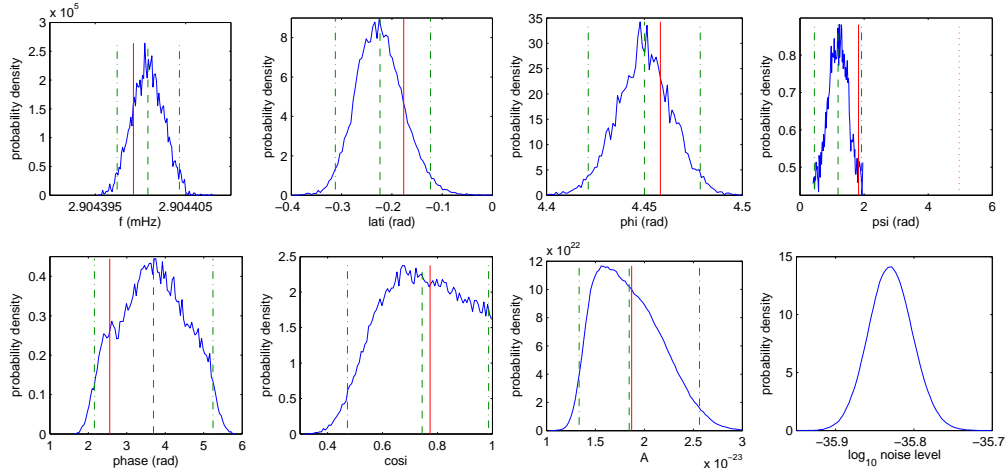


Figure 3. The same plots as in Figure 2, but for Challenge 1B.1.1b data set.

stage of five different runs, each of which used different initial data and seeds for the random number generator.

In Figures 2 and 3, and Table 2, we summarise the results of the analysis. In order to assess the performance of our approach we compare our results with the parameters that describe the actual signal present in the data set – the “key” file – and also the outcome of the analysis from other participating groups. Our analysis pipeline performed in general in a satisfactory way: the true parameters of the signals are all within the 95.5% posterior probability interval, and the difference between the true and recovered mean value of each parameter is comparable with that obtained by other groups. Notice that the accuracy of the measurements varies considerably for the seven parameters of the signal: the two sky location angles can be estimated with an error that is typically of the order of 0.1 rad, but the error on the inclination angle is around 0.5 rad, and for the polarisation angle and initial phase it increases to ≈ 0.7 rad and ≈ 1.5 rad, respectively. The signal frequency can be estimated with very high accuracy, 3.5 nHz, which represents around one tenth of a frequency resolution bin of width 32 nHz. Finally, the logarithm of the noise level estimation provided by the MCMC code has a typical error of 0.08, whereas the relative error of the signal’s amplitude estimation is typically 35%.

Additional quality indicators of the results of a given analysis that are used the context of the MLDCs, are the *correlation* \mathcal{C} between the true waveform h_{key} and the recovered one h_{rec} , and the recovered signal-to-noise ratio SNR_{rec} with respect to the optimal one SNR_{key} [23]. Our results yield \mathcal{C} that differs from 1 by less than 1%, and $\text{SNR}_{\text{rec}} = 13.577$ ($\text{SNR}_{\text{key}} = 13.819$) for Challenge 1B.1.1a and $\text{SNR}_{\text{rec}} = 23.479$ ($\text{SNR}_{\text{key}} = 24.629$) for Challenge 1B.1.1b. Both figures of merit are satisfactory and comparable with those obtained by other participating groups, see Table 2, and Table 1 of [23].

In summary, the outcome of the analysis on the Challenge 1B.1.1a-b data sets provides a successful validation of our analysis pipeline for the simplest possible scenario, and suggests that the inner core of our approach is sound and can be regarded as a solid starting point for more complex searches. However, a realistic analysis of LISA data needs to deal with tens of thousands of white dwarf binary

Table 2. Summary of the results for Challenge 1B.1.1a-b. We show the mean values and the 95.5% probability interval of the recovered parameters. We also present the true values from key files, the difference between them and our submitted results [$\Delta\lambda$ (ours)] and the median of the results submitted by other participating groups [23] [$\Delta\lambda$ (typ)]. We also show the *optimal* SNR (from the key file), the *recovered* SNR (values of SNR in $\Delta\lambda$ columns represent $|\text{SNR}_{\text{key}} - \text{SNR}_{\text{rec}}|/\text{SNR}_{\text{key}}$), as well as the correlation between the true waveform and the recovered one, \mathcal{C} .

Challenge 1B.1.1a					
Parameter	Mean and 95.5% int.		True	$\Delta\lambda$ (ours)	$\Delta\lambda$ (typ)
f (mHz)	1.0599865	$+3.5 \times 10^{-6}$ -3.9×10^{-6}	1.0599880	1.51×10^{-6}	1.81×10^{-6}
ℓ (rad)	-0.57	$+0.13$ -0.14	-0.575	0.005	0.017
ϕ (rad)	6.199	$+0.162$ -0.148	6.278	0.079	0.072
ψ (rad)	1.39	$+0.73$ -0.74	3.824	0.708	0.700
φ_0 (rad)	5.81	$+1.56$ -1.52	0.850	1.324	1.320
ι (rad)	0.81	$+0.39$ -0.63	0.637	0.173	0.248
$\mathcal{A} \times 10^{23}$	5.99	$+3.15$ -2.16	5.34	0.647	1.220
SNR	13.577		13.819	0.018	0.027
\mathcal{C}	ours: 0.992			typical: 0.988	

Challenge 1B.1.1b					
Parameter	Mean and 95.5% int.		True	$\Delta\lambda$ (ours)	$\Delta\lambda$ (typ)
f (mHz)	2.904401	$+3.0 \times 10^{-6}$ -4.0×10^{-6}	2.9043992	1.78×10^{-6}	1.92×10^{-6}
ℓ (rad)	-0.22	$+0.10$ -0.09	-0.176	0.044	0.061
ϕ (rad)	4.45	$+0.029$ -0.029	4.458	0.008	0.009
ψ (rad)	1.19	$+0.73$ -0.75	4.968	0.636	0.736
φ_0 (rad)	3.69	$+1.55$ -1.53	2.556	1.134	0.987
ι (rad)	0.73	$+0.35$ -0.56	0.687	0.043	0.123
$\mathcal{A} \times 10^{23}$	1.84	$+0.72$ -0.51	1.87	0.029	0.048
SNR	23.479		24.629	0.047	0.060
\mathcal{C}	ours: 0.996			typical: 0.981	

systems, distributed over a wide (\sim tens of mHz) frequency window and for a range of signal-to-noise ratios.

Due to the nature of the A parameter, we should expect the distribution of $\log A$ to be approximately symmetrical, whereas we have sampled and shown A itself, which is asymmetrical. The $\cos(\iota)$ parameter is correlated with A , and so its posterior distribution is also asymmetrical.

At the time of the deadline for Challenge 1B our analysis approach was still immature and the first natural extension of the code – *i.e.* the ability to search automatically over a wider frequency band than that required for Challenge 1B.1.1a-b – was still under development. In fact, results were not submitted for Challenge

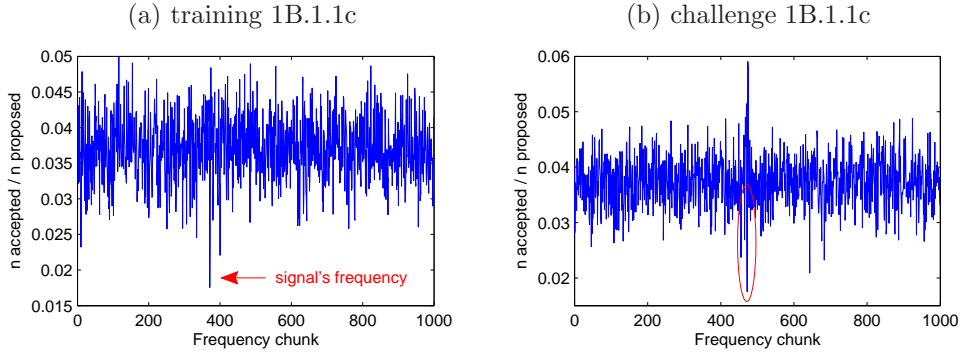


Figure 4. Looking for the frequency value in Challenge 1B.1.1c, we used the final ratio between the number of accepted transitions in f , over the number of proposed, as a signal presence indicator, expecting to find a minimum. Here we plot this quantity for the 1000 different frequency chunks we divided the data into, for (a) the training data set and (b) the challenge data set. For this latter case, we highlight with an ellipse the absolute minimum of this quantity, expecting to find a signal in that frequency region.

1B.1.1c. For completeness, in the next Section we shortly summarise the problems that were encountered and justify the lack of an entry for Challenge 1B.1.1c. This also provides a natural justification for the on-going development work.

3.2. Challenge 1B.1.1c

The efficiency of an MCMC search method for galactic binaries depends crucially on the width of the prior frequency range of the signal. For Challenge 1B.1.1a-b we were able to further restrict the prior range following the *ad hoc* strategy described in the previous Section. This was not possible for Challenge 1B.1.1c, as no clear peak(s) could be identified in the raw power spectrum of the data. This is due to several (simple) reasons: the wider (by a factor 10) prior frequency interval and the higher frequency of the signal, that produces a larger absolute Doppler frequency shift and therefore spreads the signal power over many frequency bins.

In principle one could just let the MCMC code run over a 2 mHz band and eventually the algorithm should converge; however this can take an unreasonably long time and does not provide any real advance in developing an algorithm for more complex and realistic situations. The interim solution that we adopted was to divide the whole frequency band into 1000 overlapping intervals of 63 frequency bins per interval, and we run the MCMC code in each of them. As an indicator for the presence of a signal, we used the final ratio between the number of accepted and proposed transitions in f , expecting to find a minimum – the Markov chain has converged and the transition acceptance probability diminishes – when a galactic binary signal is indeed present.

We tested this procedure on the training data set and we found a minimum in the frequency chunk that contained the signal, see Figure 4a. When running the analysis on the challenge data set, the results contained some additional complication, and this is reported in Figures 4b and 5. We could correctly establish that the signal was confined to the region spanned by two frequency chunks (corresponding to the number 471 and 472, see Figure 4b) covering the band 9.942 mHz–9.946 mHz, which is

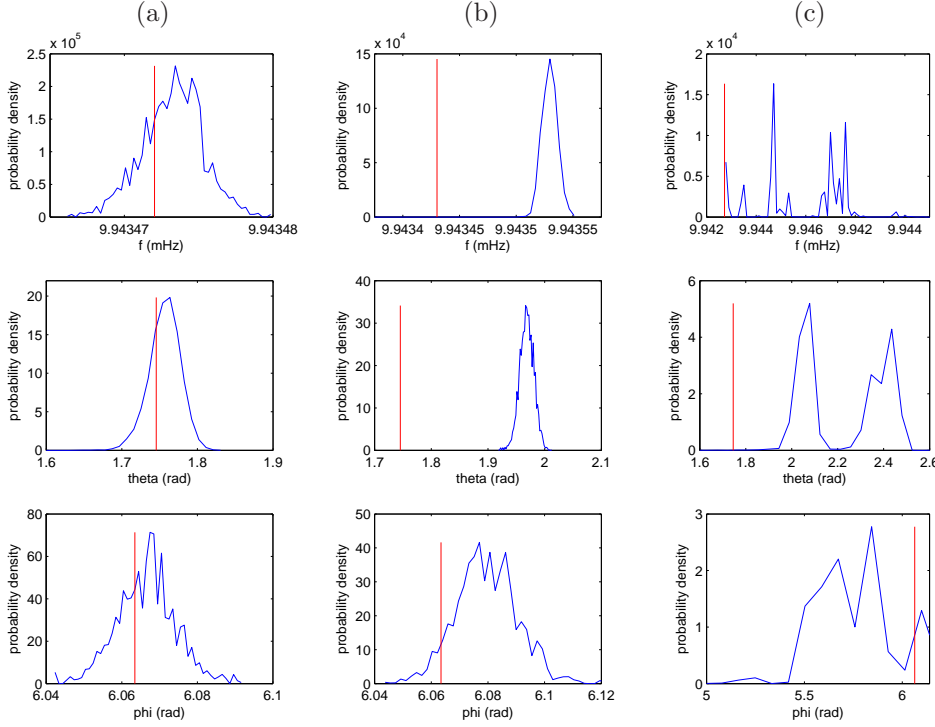


Figure 5. Marginalised PDFs of frequency and sky location (here, θ represents the co-latitude) parameters, after running our MCMC code on the Challenge 1B.1.1c data set around the frequency value identified from Figure 4b. Cases (a), (b) and (c) correspond to independent results obtained just by changing the starting values of the chain and the initialisation of the random number generator.

consistent with the true value of the key file, $f_{\text{true}} = 9.94347$ mHz. However, running three independent chains with arbitrary starting values and different initialisation seeds produced posterior PDFs which provided ambiguous results. In Figure 5 we plot the posterior PDFs of three parameters (frequency and sky location), for the three chains. We see that for case (a) the MCMC converged on the correct point in the parameter space, but for (b) the chain converged onto a secondary maximum of the likelihood function, and in (c) it did not even converge. This is the effect of a known property of the likelihood function produced by galactic binary signals that is characterised by multiple peaks at a frequency separation $\approx 1 \text{ yr}^{-1}$ [5] that our code was not yet sufficiently mature to deal with. Some *ad-hoc* tests could be used in order to discriminate the chains that converged to the actual place in the parameter space from the other ones, like comparing the recovered SNR or looking at the residuals when the recovered source was subtracted, but we are focussing on the future applications of the algorithm where it will be necessary to have all the chains sampling properly the posterior PDFs in a reasonable number of steps. Thus, work is currently on-going to address this issue, while retaining the Markovian properties of the chains. It should be noted that the problem of multimodal posterior distributions is more general, and also present in other sources considered so far in the MLDCs, for example massive-black-hole binary inspirals and extreme mass-ratio inspirals.

4. Conclusions and future work

We have begun the development of an end-to-end analysis algorithm to identify and study stellar-mass galactic binary systems in a Bayesian framework. The core of the analysis is based on an application of MCMC techniques. Using a first implementation of some of the building blocks of this pipeline we have successfully analysed the simplest single-source Challenge 1B.1.1a-b data sets: the true parameters of the signals all lay within the 95.5% posterior probability interval, the combined A and E recovered signal-to-noise ratio exceeds 95% of the optimal one and the correlation between the recovered and the true waveform is above 99%. This performance is competitive with that achieved by other groups who analysed the same data sets. Due to the immaturity of the algorithm in dealing with the multiply-peaked structure of the likelihood function for galactic binary signals at higher frequencies, results were not submitted for the Challenge 1B.1.1c data set.

Work is currently on-going to address the limitations of the present MCMC implementation. In particular, an extension to this algorithm using a Delayed Rejection MCMC technique has been developed and is undergoing testing and validation. A description of this extended technique is in preparation. Our long-term goal is to implement a multi-source version using RJMCMC in order to perform model-selection on the number of sources.

Acknowledgments

We are grateful to Alexander Stroeer for sharing his MCMC code [11], which facilitated the development of the software used in this work. MT acknowledges the University of Birmingham for hospitality while this work was carried out and is grateful for the support of the Spanish Ministerio de Educación y Ciencia Research Projects FPA-2007-60220, HA2007-0042, CSD207-00042 and the Govern de les Illes Balears, Conselleria d'Economia, Hisenda i Innovació. AV and JV acknowledge the support by the UK Science and Technology Facilities Council.

References

- [1] Bender P, Danzmann K and the LISA Study Team 1998 Laser Interferometer Space Antenna for the Detection of Gravitational Waves, Pre-Phase A Report **MPQ 233** (Garching: Max-Planck-Institut für Quantenoptik)
- [2] Nelemans G, Yungelson L R and Portegies Zwart S F 2001 *Astron. Astrophys.* **375**, 890
- [3] Farmer A J and Phinney E S 2003 *Mon. Not. Roy. Astron. Soc.* **346**, 1197
- [4] Timpano S E, Rubbo L J and Cornish N J 2006 *Phys. Rev. D* **73** 122001
- [5] Cornish N J and Crowder J 2005 *Phys. Rev. D* **72**, 043005
- [6] Crowder J, Cornish N J and Reddinger L 2006 *Phys. Rev. D* **73**, 063011
- [7] Crowder J and Cornish N J 2007 *Phys. Rev. D* **75**, 043008
- [8] Crowder J and Cornish N J 2007 *Class. Quant. Grav.* **24** S575
- [9] Stroeer A *et al.* 2007 *Class. Quant. Grav.* **24** S541
- [10] Umstätter R *et al.* 2005 *Phys. Rev. D* **72**, 022001
- [11] Stroeer A, Gair J and Vecchio A 2006 *AIP Conf. Proc.* **873**, 444
- [12] Cornish N J and Littenberg T B 2007 *Phys. Rev. D* **76** 083006
- [13] Cornish N J and Porter E K 2007 *Phys. Rev. D* **75** 021301
- [14] Cornish N J and Porter E K 2007 *Class. Quant. Grav.* **24** S501
- [15] Rover C *et al.* 2007 *Class. Quant. Grav.* **24** S521
- [16] Brown D A, Crowder J, Cutler C, Mandel I and Vallisneri M 2007 *Class. Quant. Grav.* **24** S595
- [17] Babak S 2008 Building a stochastic template bank for detecting massive black hole binaries *Preprint* arXiv:0801.4070 [gr-qc]

- [18] Harry I, Fairhurst S and Sathyaprakash B S 2008 A Search for Super Massive Black Hole Coalescences in the Mock LISA Data Challenges, in this volume
- [19] Gair J, Babak S, Porter E K and Barack L 2008 A constraint Metropolis-Hastings search for EMRIs in the Mock LISA data challenge round 1B, in this volume
- [20] Cornish N J 2008 Detection Strategies for Extreme Mass Ratio Inspirals, in this volume
- [21] Arnaud K A *et al.* 2006 *AIP Conf. Proc.* **873**, 619
- [22] Arnaud K A *et al.* [Mock LISA Data Challenge Task Force Collaboration] 2006 *AIP Conf. Proc.* **873**, 625
- [23] Babak S *et al.* 2008 The Mock LISA Data Challenges: from Challenge 1B to Challenge 3, in this volume
- [24] Mock LISA Data Challenge Task Force home page: <http://astrogravs.nasa.gov/docs/mldc>
- [25] Prix R and Whelan J T 2007 *Class. Quant. Grav.* **24** S565
- [26] Whelan J T, Prix R and Khurana D 2008 Mock LISA Data Challenge 1B: improved search for galactic white dwarf binaries using an F-statistic template bank, in this volume (*Preprint* arXiv:0805.1972)
- [27] Jaranowski P, Królak A and Schutz B F 1998 *Phys. Rev. D* **58** 063001
- [28] Gamerman D 1997 *Markov Chain Monte Carlo: Stochastic Simulation of Bayesian Inference* (London: Chapman & Hall)
- [29] Prince T A, Tinto M, Larson S L and Armstrong J W 2002 *Phys. Rev. D* **66** 122002
- [30] The Mock LISA Data Challenge Task Force Collaboration 2007 MLDC Round 2: Guidelines for returning results, from <http://astrogravs.nasa.gov/docs/mldc>